

The gender pay gap in the UK: children and experience in work

Monica Costa Dias*, Robert Joyce†and Francesca Parodi‡

February 2018

Abstract

Despite some convergence, the gender pay gap remains large. In this study, we use BHPS-USoc data to document the evolution of the gender pay gap in the UK over the past 25 years and its association with fertility. We also investigate the potential role of various differences in career patterns between men and women and how they change with the arrival of the first child. We show that differences in accumulated years of experience and in working hours play an important role. We develop an empirical wage model to estimate the causal effect of working experience in the wages of women. Estimates from this model are then used to simulate counterfactual scenarios where women always work full-time if at all and where women work as much as men do. We find that differences in working experience can explain up to two thirds of the existing gender pay gap of college graduates 20 years after the first childbirth, and that the gap is largely driven by differences in working hours. The role of working experience is more moderate for individuals with no college education, but it can still account for about one third of the overall gender wage gap 20 years after childbirth.

1 Introduction

Gender wage differentials remain substantial and reducing these differences is high on the political agenda. Understanding these differences is important not only from the point of view of gender equality per se, but also for how best to address low pay and a lack of wage progression more generally. Poverty is increasingly a problem of low pay rather than lack of employment. The proportion of people in paid work has reached record levels, with female employment having risen especially quickly over the last 25 years, and two-thirds of children in poverty now

*Institute for Fiscal Studies and University of Porto. monica_d@ifs.org.uk

†Institute for Fiscal Studies. robert_j@ifs.org.uk

‡University College London and Institute for Fiscal Studies. francesca.parodi.13@ucl.ac.uk

live in a household with someone in paid work. Understanding the wage gap between men and women is important in its own right, but all the more so now that so many families are left in poverty as a result of low wages.

In principle, there are many reasons why the wages of male and female workers might be different. To name a few of the possibilities, they could have different levels of education or labour market experience; they could be in different kinds of jobs offering different balances between financial benefits (such as wages) on the one hand and other benefits (such as flexibility in hours) on the other; they could be working in different local labour markets, with different degrees of competition for workers between employers, putting different amounts of upward pressure on wages; they could bargain differently over their wages; or there could be outright discrimination. In this paper we document the recent trends in the gender pay gap in the UK, we describe how it evolves over the course of life and its association with fertility, and we explore the power of many different alternative drivers of the gap and its evolution over the course of life, including education, working experience, working hours, occupation, industry and a range of job characteristics. We then focus on working experience and develop a simple but flexible model that we use to estimate the causal effect of cumulative years of work and working hours on the wages of women. We use our estimates to quantify the role of differences in labour supply in driving the gender pay gap in the UK.

Some of the channels we investigate here have been studied in the literature. For instance, Card et al, 2016, focus on sorting and bargaining as drivers of gender pay differences; Adda et al, 2017, show that losses in skills associated with career interruptions and sorting are important drivers of the gender gap and relate to fertility decisions; Bertrand et al. 2010, identify career interruptions and working hours as key drivers of the pay gap among MBA graduates; Goldin, 2014, identifies the high penalty for career breaks and flexibility in some high-wage occupations as a mechanism behind the persisting gender wage gap; Blau and Kahn (2017) find that the three most important factors explaining the current gender wage gap in the US are occupation, industry and experience.¹

Our study adds to this literature by measuring the role of differences in labour supply choices, including both in employment and working hours, in driving the gender pay gap for a representative sample of all workers in the UK. We also document the role of other possible drivers of the pay gap in the UK and show that their impact is likely to be more modest than that of working experience.

We find that accumulated working experience is an especially important driver of differences in pay between college graduate men and women. For them, we estimate that the divergence in labour supply that happens after the birth of the first child amounts to about two thirds

¹Earlier research focused on the role of discrimination and pre-market factors, such as education and family background, see Altonji and Blank, 1999.

of the gender pay gap 20 years later. In contrast, accumulated working experience can only explain about one third of the gender pay gap of workers who do not have a college degree. In all cases, working part-time hours seems to be key in explaining pay differences. We estimate that part-time working hours are linked to no wage gains, and this result seems to be key in explaining the wage stagnation that women face once they become mothers.

The rest of the paper is organized as follows. Section 2 describes the data and defines the main variables. Section 3 presents some descriptive evidence on the evolution of the gender pay gap, its relation to childbirth and to the the changing individual and job characteristics over the course of life. Section 4 outlines the econometric model and estimation technique we use to identify the causal impact of working experience on the wages of women. Section 5 reports the empirical estimates and 6 discusses the model predictions by simulating counterfactual wage gaps if women were to work at the same rate as men do. Finally, section 7 draws some concluding remarks.

2 Data

We use two datasets in our empirical analysis, the British Household Panel Survey (BHPS) and the Understanding Society (USoc).² The two datasets have a very similar structure and cover a similar set of variables but for different time periods. Indeed, USoc was designed to replace the BHPS and extend its sample, and the two datasets have been harmonised. BHPS covers the 1991 to 2008 period and USoc starts in 2010 and now covers the period up to 2015. We call the merged data BHPS-USoc.

BHPS-USoc is a longitudinal annual survey of families. All individuals in the original 1991 sample and subsequent booster samples remain in the panel from then onwards, apart from some lost because of attrition. This is true even in the transition from the BHPS to the USoc, although the attrition was particularly strong at that point with about 31% of the BHPS sample in 2008 being lost. USoc, however, has a substantially larger sample than the BHPS as many new families were added in 2010. Other individuals have been added to the sample in different periods - sometimes temporarily as they formed families with original interviewees or were born into them. All members of the household aged 16 and above are interviewed yearly, and a large set of demographic, educational, and labor-market information is recorded, including historical data on past working spells, working hours and socio-economic background.

We focus on the main working years between ages 20 and 55, and consider the sub-sample of men and women in this age interval who have finished education. In each year, we observe

²University of Essex. Institute for Social and Economic Research, NatCen Social Research, Kantar Public. (2017). Understanding Society: Waves 1-7, 2009-2016 and Harmonised BHPS: Waves 1-18, 1991-2009. [data collection]. 9th Edition. UK Data Service. SN: 6614.

employment status, usual working hours as well as paid and unpaid overtime, gross pay, including for overtime. Our summary measure of hours of work for those in work is obtained by dividing hours into part-time and full-time, corresponding to 5-24 and 25 or more hours of work per week respectively. Individuals reporting 4 or less weekly hours of work are considered to be not working.

Table 1 shows the sample sizes in BHPS-USoc for our population of interest. There are about 12,000 BHPS respondents, 50% of whom are observed for 5 periods or more. USoc is larger, with almost 24,000 respondents but observed for a shorter period given the time span of the data. In USoc, 50% of the sample is observed for 2 or more periods. Over 20% of the BHPS sample is followed into the USoc period, and for them we have comparatively observation windows of 11 or more years in 50% of the cases.

Table 1: BHPS-USoc – sample sizes and distribution of education

	Men	Women	All
Sample size: number of individuals			
BHPS	5,667	6,667	12,334
USoc	11,346	12,563	23,909
BHPS and USoc	1,279	1,450	2,729
Sample size: number of observations			
BHPS	37,241	47,017	84,258
USoc	41,538	45,529	87,067
BHPS and USoc	14,484	16,755	31,239
Median duration of observation spells (years)			
BHPS	5	6	5
USoc	2	2	2
BHPS and USoc	11	11	11
Distribution of education			
GCSEs	.456	.438	.446
A-levels	.406	.434	.421
Degree	.137	.126	.131

We consider three education groups: GCSEs, representing those who leave education at 16 without completing high school education; A-levels, representing those with a high-school diploma or equivalent; and Degree, representing those who graduate from college (3-year degree). The distribution of education in the population is shown at the bottom of Table 1. The largest group is that with lower education attainment, and this is true for both men and women. Men in our age group and time window are more likely to have a Degree than women, but as we will show, this does not hold in the more recent periods. Only about 13% of our sample has a degree.

Our measure of the hourly wage rates is the ratio of the total gross weekly pay by total hours, including paid and unpaid overtime. We then remove aggregate wage growth from the wage rate and trim it at percentiles 2 and 99 from below and above, respectively, to limit the impact of measurement error in wages and hours. All our results are for wages are in 2015 wage levels.

The historical labour-market information is collected in some waves only and only for those who join the sample since the last collection of historical data. We have historical information for about 60% of our sample. It allows us to construct two experience variables that measure, respectively, actual accumulated experience time in part-time and full-time work since the beginning of the working life. We then complete the experience variables over the entire observation window using year-on-year information on employment spells and hours.

Individuals who are working at the time of the interview are asked to report their occupation and industry, and these are classified using standard classification codes, SOC and SIC respectively. During the span of our data, there were changes to these classification codes. For comparability over time, we convert all classifications into the most recent ones, SOC2010 (3-digit) and SIC2007 (2-digit). As there is no exact one-to-one mapping between subsequent versions of these classifications, we construct gender specific conversion matrices, whose rows contain the (cumulative) probabilities with which each code of the former classification is translated into different codes of the new classification. We then use the conversion matrix to assign randomly a code in the new classification to observations collected under the old classification.


The British Labour Force Survey (LFS) contains more detailed information about the characteristics of jobs than the BHPS and relies on a much larger sample. We use it to gather information about the ‘women/family friendliness’ of jobs by industry and occupation that we then merge with the BHPS data. Specifically, we merge in information on the percentage of employees who are women by occupation, percentage of workers doing part-time hours by occupation and percentage of managers who are women by industry.

3 Descriptive analysis

3.1 Differences in wages between men and women over time

We start by looking at the long-term trends of the gender wage gap in the UK. As for many other developed economies, gender wage disparities in the UK remain high despite a steady convergence over time since the 1980s. Figure 1 plots the average hourly wages of male and female employees over time according to the LFS. LFS data is used to produce the time trends presented in this section because its larger sample size capture the trends more accurately; results obtained using BHPS-USoc data produce similar but more irregular patterns. It also plots, in black and on the right hand axis, the proportional difference between the two. The gap has decreased over the last twenty years from almost 30% in 1993. Currently, the average female employee earns around a fifth less per hour than the average male employee.

Figure 1: real hourly wages of men and women over time




Notes: LFS. Wage rates per hour in constant 2016 wage levels; observations in the top one and bottom two percentiles of the wage distribution by gender and year are excluded. Wage gap measured in proportion to male wages.

This wage gap is what it says on the tin: the difference between average female wages and average male wages. It is not a like-for-like comparison between otherwise-identical workers or jobs. One reason why wage differentials between men and women might have changed over this period is that their relative levels of education have also change. This is actually an important

aspect to take into account in interpreting the declining wage gap over time. Figure 2 shows that the population has become more highly educated at a rapid rate over the past 20 years, with a rapid rise in the proportion of graduates and a rapid fall in the proportion of people with no more than GCSE-level qualifications. It also shows that women have experienced the more rapid increase in education levels. In fact, in the late 2000s, they 'overtook' men in this respect: women are now, on average, more highly educated than men. Because graduates tend to earn more than non-graduates these differential trends in educational attainment have contributed to reducing the gender wage gap.


Figure 2: Educational attainment for men and women over time



Source: LFS.

Figure 3 shows the evolution of the gender wage gap as a proportion of male earnings over time, by education. For those with GCSE-level qualifications, this plot confirms that indeed the gender pay gap has fallen over the past two decades. However, it reveals no clear downward trend for the other education groups. As a result there has been a notable change in the nature of the gender wage gap. The gap used to be bigger (in proportional terms) for the least well educated women than for graduates; whereas the reverse is now the case. In summary, the fall in the overall gender wage gap over the past 20 years has been driven mostly by the lowest-educated individuals, and by an increase in the number of women who are highly educated.

Figure 3: Gender wage gap by education level over time




Notes: LFS. Wage gap measured in proportion to male wages.

3.2 Children, career patterns and the gender wage gap over the lifecycle

A crucial starting point for disentangling the drivers of wage differences between men and women, which simple aggregate figures miss, is that those differences evolve over the life cycle. This in turn is highly related to the arrival of children and changes in labour market behaviour associated with that. Figure 4 shows how average wages for male and female employees relate to their age (pooling those observed at the relevant ages between the start of 1993 and the middle of 2017). Notice that the sets of individuals who are employed at each age are different, so it is possible, for example, that women with low levels of experience return to employment in their 40s, thereby dragging down average female wages at that age. Wages are shown in 2016 constant-wage terms, so the increasing profile with age means that wages increase over the course of life by more than would be expected simply due to economy-wide growth. The figure shows that wages typically increase with age throughout the 20s, for both men and women, which is consistent with the returns to additional experience being especially high for those with little experience. These returns look higher for graduates: their wage profile is especially

steep throughout their 20s and, for men, well beyond that.

Figure 4: Mean hourly wages across the life cycle by gender and education




Notes: LFS 1993 to 2017. Wage rates per hour in constant 2016 wage levels; observations in the top one and bottom two percentiles of the wage distribution by gender and year are excluded.

The gender wage gap is small or non-existent at around the time of labour market entry and it widens only slowly up to the mid 20s, particularly for college graduates. The gap then opens up more from around the late 20s and gets gradually wider over the next 20 years of the life cycle. This is because male wages continue to increase, especially for the high-educated, while female wages completely flatline on average.

The opening of the gender wage gap when people reach their late 20s is likely related to the arrival of children. Figure 5 shows this explicitly by plotting the wage gap not by age, but by time to or since the birth of the first child in a family (where zero is the year in which that child is born). There is, on average, a wage gap of over 10% even before the arrival of the first child. A small part of this gap is simply due to age differences - men tend to be slightly older than women when the first child arrives though the age-adjusted line on Figure 5 still yields a wage gap of 7-12% in the five years preceding the first child. A key feature of the patterns shown in the figure is that the gap appears fairly stable until the child arrives, and is small relative to what follows: after the child arrives, there is a gradual but continual rise in the wage gap over the following 12 years, until it reaches around one third.

Figure 5: Gender wage gap by time to/since birth of first child




Notes: BHPS-Usoec 1991 to 2015. Wage gap measured in proportion of male wages. Observations in the top one and bottom two percentiles of the wage distribution are excluded. The age-adjusted series shows the gap obtained from wage rates net of education-specific age effects.

While we measure wages on an hourly basis, and hence differences in working hours cannot directly explain the gender pay gap described above, different working patterns may lead to different hourly pay for more subtle reasons related to productivity, career progression or other job characteristics. Figure 5 suggests that changes in women’s working patterns after the arrival of children may well be important in explaining this wage gap. The crucial observation is that the wage gap opens up gradually – not in any sudden jump – after the first child arrives and continues to widen for many years after that point. This pattern would be consistent with a gender gap in the level of labour market experience following the same basic shape as the gender gap in wages: relatively stable in the years before childbirth, growing incrementally for many years after that point, before eventually largely stabilising once more. The next 3 figures show this.

Figure 6 shows the employment rates of men and women by the time to or since the birth of the first child. Before the arrival of the first child, it is difficult to discern any differences between the employment rates of men and women. However, when the first child arrives, a large employment gap opens up immediately: many women leave paid employment at this point, while any employment responses by men look tiny in comparison, and non-existent

for high-school and college graduates.³ The employment response among the lowest-educated women is more than double the response among other women. Between one year before and one year after the birth of the child, women's employment rates drop by 30 percentage points (ppts) for those with GCSEs, 13ppts for those with A levels and 9ppts for graduates. The other important feature of Figure 6 is that, once the employment gap opens up after the arrival of the first child, it persists. Women's employment rates do start to rise again once the first child is around school age, but they remain below male employment rates for the full 20 years shown. Hence, the gap in time spent in paid work keeps growing year on year for a long time after the first child arrives.

Figure 6: Employment rates of men and women




Notes: BHPS-UsoC 1991 to 2015.

Figure 7 shows that not only do many women move out of paid employment altogether after having their first child, but many others move to work that is part-time (recall that part-time is defined as working 5 to 24 hours per week). Again, the male rate of part-time employment looks essentially unaffected by the arrival of the first child, and the gap that opens up is persistent: women are still significantly more likely to be in this kind of work than men when their first child reaches adulthood.

³For the purpose of measuring employment, maternity or paternity leave is treated as being in paid work.


Figure 7: Proportion of all men and women in jobs of no more than 25 hours per week



Notes: BHPS-UsoC 1991 to 2015.

Figure 8 shows the direct implications of these patterns: a steadily increasing gap in accumulated labour market experience after the arrival of the first child. By the time their first child is aged 20, women have on average been in paid work for three years less than men, comprising ten years less paid full-time work and seven years more part-time paid work. The gap is larger still for the low-educated. Previous research (Blundell et al., 2016) tells us that the three years less spent in any form of paid work understates the gender differences in accumulated human capital it is the ten-year gap in full-time experience that is more relevant. This is because it is only full-time paid work which seems to have substantial benefits in terms of the accumulation of experience that allows workers to command higher wages in future. We confirm this in the new analysis summarised in the next section, and examine the implications of the lack of wage progression in part-time work for the gender wage gap.

Figure 8: Gender gap in years spent working full-time and part-time




Notes: BHPS-Usoc 1991 to 2015. The figure cumulates the gender gaps in years of work shown in Figures 6 and 7, and therefore does not include any differences in experience that already exist more than five years before the birth of the first child; these earlier differences are negligible.

There are certainly other factors, besides levels of experience in paid work, that may be affected by childbirth and that may contribute to differences in wages between men and women. One possibility is that women undertake different kinds of work upon becoming mothers, potentially in different sectors of the economy. Such change in job characteristics could be related to their wages for a number of reasons. For example, priorities or constraints could change around the time that children arrive such that women move towards occupations in which the benefits are less skewed towards wages and more towards other factors, such as flexibility. It could also be that a concentration of women in certain occupations or industries allows employers to exercise market power in order to hold wages down if, for example, they know that many of those women have limited ability or desire to search for alternative employment because they are time-constrained or want to work close to home. These different kinds of mechanisms linking occupation, industry or other job characteristics to the gender wage gap would have very different implications for policy, and it is beyond the scope of this work to disentangle them (and there are many other possibilities besides the examples given). But what we can do is provide a sense of their likely importance in accounting for the evolution of the gender wage gap.

Figure 9 summarises three example differences between the occupations and industries that women and men work in, and how these differences evolve at around the time of childbirth. We take the occupation or industry that each worker is in, and map this to the composition of the workforce in that occupation or industry (computed from the LFS). As time goes on, women who have children tend (relative to men who have children) to concentrate increasingly in female-dominated occupations, occupations in which part-time work is relatively common, and sectors in which female managers are relatively common. To that extent there are similarities with the evolution of the gender wage gap - which also grows over the lifecycle, as we have

seen. However, a closer look reveals a caveat to that: whereas the gender wage gap is fairly stable in the years before childbirth and then begins gradually increasing from the time of the first child, occupation and industry differences between men and women seem to be on a more uniformly increasing trajectory that starts a few years before the birth of the first child. This may in part be due to job changes in anticipation of having children. But it casts some doubt on the ability of these occupation or industry differences to explain powerfully the shape of the gender wage gap over the lifecycle documented above.

Figure 9: Gender gaps (women minus men) in characteristics of occupation and industry



Notes: BHPS-UsoC 1991 to 2015.

So far we have highlighted numerous factors that can play a role in driving the gender wage gap that persists in the UK labour market: education, labour supply choices along the intensive and extensive margins and the resulting work experience accumulation patterns, characteristics of specific occupations/sectors, jobs and working arrangements. We now perform a more comprehensive decomposition exercise that allows us to quantify the association between the gender wage gap and these factors. To do so, we estimate a set of wage regressions. We start from a baseline specification that only accounts for demographic controls (age and region). The gender pay gap conditional on this variables is captured by the coefficient of a female dummy. We then proceed to richer specifications by progressively controlling for a number of differences between male and female workers: education and education interacted with age, part-time and full-

time experience polynomials interacted with education, industry (2-digit SIC2007 codes) and occupation (3-digit SOC2010 codes), other characteristics of the job including working hours, indicator for public sector and manual jobs, flexible working arrangements, share of co-workers who are women within occupation, share of co-workers doing part-time hours within occupation, shares of co-workers by education within occupation and share of women in managing positions within industry.]


Table 2 and Figure 10 displays the estimate coefficients for the female dummy and how it varies as more detailed information about individual characteristics, working history and the characteristics of the jobs is added to the regression model. The first row in the table shows that the raw gap of 22% is only mildly reduced by accounting for gender differences in educational attainment. Gender differences in experience show by far the strongest impact in reducing the gap (compare the estimate in column 3 to those in columns 1 and 2). After controlling for experience in full-time and part-time jobs, the gender gap in wages drops to just below 10%. After that, differences in industry and occupation can further reduce the gap by another 2ppts (column 4), and other job characteristics have not further impact (column 5).

Table 2: log hourly wage regressions, all education levels

	(1)	(2)	(3)	(4)	(5)
female	-0.224*** (0.003)	-0.195*** (0.003)	-0.086*** (0.003)	-0.069*** (0.003)	-0.071*** (0.003)
age, region	yes	yes	yes	yes	yes
education	no	yes	yes	yes	yes
experience	no	no	yes	yes	yes
occupation, industry	no	no	no	yes	yes
job characteristics	no	no	no	no	yes
N	72313				

Figure 10 shows these results in more detail, by years to/since first childbirth. It highlights that experience has the potential to account for a large amount of the gender wage gap, including the way that it evolves over the life cycle albeit still leaving a substantial part of the gender wage gap unexplained (and our causal analysis in the next section confirms this). Industry and occupation differences, by contrast, seem to explain far less. We now turn to examine the causal role of the experience differences.

Figure 10: Gender wage gap by time to/since birth of first child, controlling for association between wages and individual and job characteristics



Note: Estimates use BHPS-Usoc 1991 to 2015.

4 Model and estimation

To measure the causal role of working experience in driving the pay differentials between men and women, we specify and estimate a simple but flexible model of experience capital accumulation and wages. In all that follows, we consider wages, experience and labour supply choices to be education-specific and, hence, model separately the choices and outcomes of different education groups. We continue to consider the same three education groups as before, corresponding to high-school dropouts, high-school graduates and university graduates (or GCSEs, A-levels and 3-year college degree in the UK education system). To simplify the notation, we omit the education index but it is implicit in all parameters below.

Following Blundell et al. 2016, wages are modeled as a simple function of accumulated experience capital. The wage rate of woman i with education level s and accumulated experience capital k_{it} at age t is

$$\ln y_{it} = \alpha_i + \ln(k_{it} + 1) + u_{it} \quad (1)$$

where y_{it} is the hourly wage rate at age t net of aggregate wage growth, α_i is an individual-

specific wage level per unit of experience capital and u_{it} is a time-varying idiosyncratic wage shock. The latter may include a persistent and a transitory component.

We assume that human capital is accumulated over time at a rate that depends on hours of work in the previous period. For those in work, we consider two hours points corresponding to part-time and full-time work. Consistently with our previous analysis, these correspond to weekly hours between 5 and 25, and 26 or more. Formally, the human capital process

$$\begin{aligned} \ln(k_{it} + 1) &= \ln(k_{it-1} + 1) - \delta + \pi(e_{it-1})P_{it-1} + \phi(e_{it-1})F_{it-1} \\ h_{i1} &= 0 \end{aligned} \quad (2)$$

where we allow for skills to depreciate in each period at a rate δ and to accumulate while in work at rates π or ϕ depending on past working hours being part-time or full-time. Moreover, the decreasing returns to additional periods of work is formalised by allowing the (π, ϕ) parameters to depend on years of work accumulated up to the start of period $t - 1$.

Our goal here is to consistently estimate the dynamic returns to work and to working hours in order to assess their combined role in driving the gender pay differentials. To do so, however, we need to deal with some important biases in estimating the wage and human capital processes in equations 1 and 2. These biases are induced by current labour supply as well as by past labour supply, in that the latter determines the accumulated working experience.

To estimate our model, we start by re-writing the wage equation (1) in first differences and replacing the growth in experience capital using equation (2). This yields

$$\Delta \ln y_{it} = -\delta + \pi(e_{it-1})P_{it-1} + \phi(e_{it-1})F_{it-1} + \Delta u_{it}.$$

We estimate this model on the sample of women working in any two consecutive periods using the control function approach advanced by Heckman (1979) to correct for three potential sources of bias. The first is the employment selection process. To control for the endogeneity of employment over two periods we add two corresponding selection terms, for current and past work respectively. The selection terms are constructed using a reduced form model for employment, which is assumed to be driven by an index $Z\beta$ where Z are the instruments for employment.

The two other sources of bias are the observed variables in the model, past working hours and years of working experience. These are endogenous in our model and potentially related to the first differences in wage shocks, Δu . In the case of accumulated years of work at $t - 1$, which results from labour supply in periods up to $t - 2$, such correlation could arise for two reasons. First, and more obviously, Δu may have a long memory and hence contain information driving past labour supply – something that would happen if, as is frequently assumed, the permanent wage shock is auto-regressive. And second, the conditioning on current and past employment may create a dependency between past working years and current wage

shocks. Such dependency could arise, for instance, in case more experienced workers are more likely to keep working when going through periods of low unobserved wage growth than less experienced workers. To address these issues, we use two reduced form models, one for each of the endogenous variables, both conditional on employment. We assume that hours of work and accumulated years of work are both driven by an index $W\gamma$ where W includes the instruments for both variables as well as those for employment.

In the empirical specification, we allow for years of full-time and part-time work to have different implications for the accumulation of skills, consistently with the human capital process in 2. Specifically

$$\begin{aligned}\pi(e_{it}) &= \pi_1 + \pi_2 e_{it}^P + \pi_3 e_{it}^F \\ \phi(e_{it}) &= \phi_1 + \phi_2 e_{it}^P + \phi_3 e_{it}^F\end{aligned}$$

where (e^P, e^F) represent past years of work part-time and full-time, respectively, and (π_j, ϕ_j) for $j = 1, 2, 3$ are unknown parameters. Our full regression model is, therefore

$$\begin{aligned}\Delta \ln y_{it} &= -\delta + (\pi_1 + \pi_2 e_{it}^P + \pi_3 e_{it}^F) P_{it-1} + (\phi_1 + \phi_2 e_{it}^P + \phi_3 e_{it}^F) F_{it-1} + \\ &\lambda_t^d(Z_{it}^d\beta) + \lambda_{t-1}^d(Z_{it-1}^d\beta) + \lambda^h(Z_{it-1}^h\gamma) + \lambda^P(Z_{it-1}^e\eta^P) + \lambda^F(Z_{it-1}^e\eta^F) + v_{it}\end{aligned}\quad (3)$$

where $(\lambda_t^d, \lambda_{t-1}^d, \lambda^h, \lambda^P, \lambda^F)$ are the control functions for, respectively, current and last period employment, last period working hours and years of work part-time and full-time hours. Equation 3 shows that the returns to a year of work for workers with no past work experience (π_1, ϕ_1) cannot be distinguished from the depreciation rate and we will, therefore, estimate

$$\tilde{\pi}_1 = \pi_1 - \delta \quad \text{and} \quad \tilde{\phi}_1 = \phi_1 - \delta.$$

We use a probit model for employment and the control function λ^d is approximated by the inverse Mills ratio.⁴ The main instrument for employment is constructed from the simulated out-of-work income – essentially a measure of the net public transfers the family is entitled to in case the woman does not work and disregarding the earnings of a present spouse. We use observed family demographics and run such information through a tax simulator – FORTAX – to predict the benefit entitlement of families had they had no earned income.⁵ We then regress the simulated disposable income variable on the variables describing family demographics that are used to calculate public transfers, and predict the residuals. These residuals vary over

⁴We have tried more flexible approximations, including polynomials of the Mills ratio and of the index, but this made no difference to the results.

⁵FORTAX is a detailed tax and benefit micro-simulation tool that can be used to accurately predict the budget constraints families face by earned income. It accounts in detail for the tax and welfare system in place at each point in time how they changed over the period of our data. More information can be found in Shephard (2009) and Shaw (2011).

time, with policy reforms, and by family demographics. It is this differential time variation by demographic group that we use as the exogenous driver of employment. We complement residual simulated income with indicators for motherhood and number of children.

In a similar vein, we use a probit to model hours of work conditional on participation and approximate the control function λ^h by the inverse Mills ratio. We use two instruments specifically targeting working hours. These are constructed from the simulated disposable income for the relevant period, for the two scenarios of women working 25 and 40 hours per week (part-time and full-time). In producing the simulated disposable income, we use predicted wages on education, age, region and a set of factors describing socio-economic background, which we describe in more detail below. We then net out the effects of family composition from the simulated incomes as before, to emphasise the differential time variation in incentives to work by demographic group that are induced by policy reforms and that we want to exploit for identification. The hours regression also controls for a set of instruments more closely related to accumulated years of working experience, including a cubic polynomial in age, a dummy for motherhood and the ages of the youngest and the oldest children.

For experience we use a linear regression model on the same set of instruments used for working hours. The control function in this case is simply the residual from this regression.

In all regression models we control for a set of exogenous matching variables that include region of residence and a quadratic polynomial on the two first factors from a principal components analysis on a group of variables describing the socio-economic background of the woman. These include her parent's education and whether they were working when she was 16, whether she lived with both parents at that same age, books at home as a child, ethnicity, number of siblings and sibling order. The background factors summarise some characteristics of the individual's parental home and are meant to capture permanent individual traits that drive productivity in the labour market and labour supply choices.

5 Estimates

We estimate our model using BHPS data for the 1991-2008 period. Estimates of the first stage regressions can be found in the Appendix to this paper, tables 4-7. They show that the instruments we are using are strong drivers of the endogenous variables in most cases. As expected, residual simulated disposable income is a stronger determinant of employment and hours of work for women with basic education only, and has less power as a driver of the same outcome among college graduates. In addition, age and age of oldest child are strong predictors of accumulated experience. We test the strength of the set of instruments meant to capture exogenous variation for each of the endogenous variables, and find evidence in support of the instruments in all cases except for hours choices among college graduates (see p-val at

the bottom of Table 5).

Using these estimates, we then construct the various control functions and include them in our regression model for wage growth. Estimates of the parameters of interest are shown in Table 3 for each of the three education levels. Columns 1, 3, and 5 display estimates by education obtained without controlling for employment selection, endogenous hours or accumulated experience; columns 2, 4, and 6 display the control function estimates. Clearly, controlling for endogenous selection and experience does not much affect the estimate of the experience effects. Figures at the bottom of the table show the p-value for the test of statistical significance of the set of control functions used to tackle endogeneity. They show marginally significant effects for the two lower education groups, but not for the top education group.

Table 3: Log wage regressions in first differences, women who left education without A-level qualifications

	GCSEs		A-levels		Degree	
	linear (1)	CF (2)	linear (3)	CF (4)	linear (5)	CF (6)
(1) Lag PT hrs	0.008 (0.012)	0.015 (0.033)	-0.034** (0.017)	0.010 (0.039)	0.028 (0.038)	-0.020 (0.079)
(2) Lag PT hrs \times lag FT exp.	-0.002 (0.001)	-0.002 (0.001)	-0.000 (0.001)	0.001 (0.002)	-0.004 (0.003)	-0.001 (0.004)
(3) Lag PT hrs \times lag PT exp.	-0.001* (0.001)	-0.001 (0.001)	0.001 (0.001)	0.001 (0.002)	-0.003 (0.003)	-0.001 (0.004)
(4) Lag FT hrs	0.048*** (0.007)	0.037*** (0.011)	0.062*** (0.007)	0.057*** (0.010)	0.075*** (0.013)	0.080*** (0.021)
(5) Lag FT hrs \times lag FT exp.	-0.002*** (0.000)	-0.002*** (0.001)	-0.003*** (0.000)	-0.003*** (0.001)	-0.004*** (0.000)	-0.003*** (0.001)
(6) Lag FT hrs \times lag PT exp.	0.000 (0.000)	0.000 (0.001)	0.000 (0.001)	-0.001 (0.001)	-0.002* (0.001)	-0.000 (0.002)
<i>F</i> -test statistic		2.442		2.364		1.198
p-val		0.087		0.038		0.309
N	7,544	7,339	6,045	5,455	3,052	2,571

Notes: Estimates on BHPS data for the 1991-2008 period. The *F* statistic is for the statistical significance of the control functions for employment, lagged employment, lagged hours and lagged experience in PT and FT hours. All regressions also include all matching variables. Bootstrapped standard errors clustered at the individual level in parentheses under the estimated coefficients.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The results in row 4 suggest that full-time work has a strong positive impact on wages, and that this effect increases with education. Moreover, the numbers in rows 5 and 6 further

suggest that the returns to one additional year of full-time work drop with full-time experience but not with part-time experience. These estimates are consistent with the view that the returns to work drop over the life-cycle as workers accumulate experience, leading to a concave wage profile over the course of working life. It is also evident from the figures in Table 3 that part-time work has little or no impact on wages. Indeed, estimates in rows 1 to 3 show small and statistically insignificant estimates of the impact of part-time work. These figures show that, on average, the wages of women working part-time hours stagnate.


6 Counterfactual simulations of the cumulative effect of employment and hours of work

We use the estimates in rows 4 and 5, columns 2, 4 and 6 of the Table to run two counterfactual experiments. First, we set all work to be full-time. We do this by setting the value of part-time work for future wages to be equal to the estimated value of full-time work. And second, we assume that women work at the same rate as men with the same education and demographic characteristics. We then simulate the experience profiles of women under these two alternative scenarios by assuming that there is no depreciation of market skills during the periods when the woman is not doing paid work. Our counterfactual wages are constructed from observed wages by netting out the experience effect at observed levels of experience and adding the experience effects at the counterfactual levels of experience. Hence, the unobserved component of wages remains unaltered in our simulations. If it changes over the life-cycle as women take different types of jobs or their wages vary for other reasons, such differential selection will still be reflected in the simulated profiles.

Figure 11 shows the results from these simulations for parents, by time since the birth of the first child and education. In each of the 3 panels, the top solid line and the bottom dash line are the observed wage profiles of fathers and mothers, respectively. The plots show how these two lines move apart as the child grows, and they also highlight that the wages of women at best stagnate after childbirth. This holds true for all education groups. In relative terms, the pay differential when the child reaches 18 is remarkably similar across education groups. In all cases, women earn about 30% less than men do at that point in the child's life.

The two intermediate lines in the graph show how employment and working hours after childbirth inhibit wage progression for women. The right hand side graph for college graduates suggest that working experience is a determinant factor in explaining the opening gap after childbirth. Given our estimates of the effect of working experience on wages, if college graduate women were to work full-time hours if they work at all, the wage gap 18 years after childbirth could close up to 50%. If, in addition, they were to work at the same rate as men do, the wage gap at the same age of the child could be further reduced to one third of the observed level.

Figure 11: Counterfactual simulations – hourly wage rate by time since first birth



Notes: Based on BHPS data for the 1991-2008 period. 2016 wage levels.

However, experience plays a less important role in determining the gender pay gap for workers who do not have a college degree. For these groups, gender differences in pay are comparatively large at childbirth and the gap in accumulated experience after that can only account for about one third of the gap in pay when the child reaches 18 years of age. Other differences in job characteristics, firm characteristics, occupation or in how wages are negotiated are likely to play a determinant role in driving gender pay differentials among workers with GCSE and A-levels qualifications.

7 Concluding remarks

Gender differences in rates of full-time and part-time paid work after childbirth are an important driver of differences in hourly wages between men and women. This is because they affect the amount and type of labour market experience that men and women build up, and this experience affects the hourly wage levels they can command. In this paper we show that differences in working experience are determinant in explaining the gender pay gap of college graduates, for whom they can explain up to two thirds of the wage differences 20 years after childbirth. The role of experience in driving the gender wage differences of those with GCSE-level and A-levels qualifications is more modest, accounting for about one third of the gap 20 years after the first

childbirth.

It is not only taking time out of paid work that matters, but crucially working part-time after childbirth seems to hold back women's wages. This is because extra experience in full-time work leads to higher hourly wages, whereas extra experience in part-time work does not.

A key challenge for future research, then, is to understand why part-time work shuts down wage progression so much. There are a number of possibilities, including less training provision, missing out on informal interactions and networking opportunities, and genuine constraints placed upon the build-up of skill by working fewer hours. Understanding this properly looks of great potential importance for policymakers who want to address the gender wage gap. Of course, our results also suggest that an alternative (or complementary) focus would be on understanding the causes of gender differences in rates of full-time work in the first place, such as the division of childcare responsibilities.

Our results also show that closing gender gaps in rates of full-time and part-time paid work, or narrowing the difference between the impacts of full-time and part-time paid work on wage progression, cannot be expected to close the gender wage gap fully. This is especially relevant when thinking about the relationship between the gender wage gap and poverty: among lower-educated people, there is already a relatively substantial gender wage gap before the first child is born, and gender differences in full-time and part-time paid work in the subsequent 20 years explain only a minority of the gender wage gap that has built up by that point. Previous research suggests that other contributing factors could include women being less likely to work in more productive firms, less likely to successfully bargain for higher wages within a given firm, and more likely to enter family-friendly occupations over high-paying ones.⁶ Better understanding of mechanisms such as these, and their underlying causes, is another key priority for further research.

References

- [1] Adda, J., C. Dustmann and K. Stevens (2017). "The career costs of children," *Journal of Political Economy*, 125: 2933-37
- [2] Altonji, J. and R. Blank (1999). "Race and Gender in the Labor Market," in Ashenfelter and Card (eds), *Handbook of Labor Economics*, vol 3C, Amsterdam: North-Holland
- [3] Bertrand, M., C. Goldin and L. F. Katz (2010). "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors," *American Economic Journal: Applied Economics*, 2(3):228-255

⁶Others have looked at some of these issues. E.g., see Adda et al 2017, Card et al. 2016.

- [4] Blau, F. D., and L. M. Kahn (2017). “The Gender Wage Gap: Extent, Trends, and Explanations,” *Journal of Economic Literature*, 55(3): 789-865
- [5] Blundell, R., Costa-Dias, M., Meghir, C. and J. Shaw (2016). “Female Labour Supply, Human Capital and Welfare Reform,” *Econometrica*, 84(5): 1705-1763
- [6] Card, D., A. Cardoso and P. Kline (2016). “Bargaining, sorting, and the gender wage gap: quantifying the impact of firms on the relative pay of women,” *Quarterly Journal of Economics*, 131: 633-86
- [7] Goldin, C. (2014). “A Grand Gender Convergence: Its Last Chapter,” *American Economic Review*, 104 (4): 1091-1119
- [8] Goldin C. and L. F. Katz (2016). “A Most Egalitarian Profession: Pharmacy and the Evolution of a Family Friendly Occupation.” *Journal of Labor Economics*, 34(3): 705-745

Appendix: Estimates of first stage regressions

Table 4: Employment regressions for women by education, probit specification

	Education		
	GCSEs	A-levels	Degree
Res. simulated inc. if $d = 0$	-0.002*** (0.000)	-0.001* (0.000)	-0.000 (0.001)
own dependent children in HH	-0.629*** (0.052)	-0.804*** (0.062)	-0.549*** (0.097)
2 children	-0.275*** (0.049)	-0.176*** (0.063)	-0.211** (0.102)
3 or more children	-0.591*** (0.072)	-0.676*** (0.094)	-0.538*** (0.165)
χ^2 test statistic	480.530	398.301	102.134
p-val	0.000	0.000	0.000
N	16,995	12,393	6,125

Notes: Estimates on BHPS data for the 1991-2008 period. The χ^2 statistic is for the statistical significance of all instruments, including simulated income and children dummies. All regressions also control for a set of background characteristics, including region of residence and a quadratic polynomial in the the two principal components from a set of socio-economic background variables (maternal and paternal education, whether lived with both parents at 16, whether parents were in paid work at that same age, ethnicity, books in parental home, number of siblings and sibling order). Clustered standard errors at the individual level in parentheses under the estimated coefficients. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Hours selection conditional on employment, women by education, probit specification

	Education		
	GCSEs	A-levels	Degree
Res. simulated inc. if $h = \text{FT}$	0.005*** (0.001)	0.004*** (0.001)	0.003* (0.002)
Diff in res. sim. inc. (FT-PT)	-0.001 (0.001)	-0.000 (0.002)	-0.004 (0.002)
age	-0.976 (1.942)	-3.763 (2.464)	-7.325* (4.029)
age squared	0.332 (0.536)	1.026 (0.690)	2.127* (1.113)
age cubic	-0.042 (0.048)	-0.098 (0.063)	-0.201** (0.101)
age youngest child	0.059*** (0.011)	0.070*** (0.016)	0.000 (0.027)
age oldest child	-0.001 (0.010)	0.008 (0.013)	0.045* (0.024)
mother	-1.443*** (0.127)	-1.794*** (0.165)	-1.211*** (0.288)
χ^2 test statistic	23.105	12.790	3.316
p-val	0.000	0.002	0.191
N	11,938	9,781	4,832

Notes: Estimates on BHPS data for the 1991-2008 period. The χ^2 statistic is for the statistical significance of the simulated income instruments in the two top rows. These measure disposable income if working full-time and the difference in disposable income in full-time and part-time, respectively. All regressions also control for a set of background characteristics, including region of residence and a quadratic polynomial in the the two principal components from a set of socio-economic background variables (maternal and paternal education, whether lived with both parents at 16, whether parents were in paid work at that same age, ethnicity, books in parental home, number of siblings and sibling order). Clustered standard errors at the individual level in parentheses under the estimated coefficients. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Years of full-time work experience among women by education, linear regression model

	Education		
	GCSEs	A-levels	Degree
Res. simulated inc. if $h = FT$	0.045*** (0.005)	0.038*** (0.006)	0.021*** (0.006)
Diff in res. sim. inc. (FT-PT)	-0.022*** (0.006)	-0.027*** (0.008)	-0.027*** (0.008)
age	13.457* (8.058)	-9.948 (8.430)	-15.112 (9.957)
age squared	-0.493 (2.324)	6.378*** (2.402)	7.170** (2.904)
age cubic	-0.120 (0.216)	-0.726*** (0.222)	-0.690** (0.277)
age youngest child	0.118** (0.059)	0.089 (0.087)	-0.297** (0.148)
age oldest child	-0.445*** (0.052)	-0.492*** (0.079)	-0.254** (0.126)
mother	2.293*** (0.566)	2.565*** (0.675)	3.386*** (0.930)
F -test statistic	159.123	181.986	114.780
p-val	0.000	0.000	0.000
N	9,481	6,845	3,212

Notes: Estimates on BHPS data for the 1991-2008 period. The F statistic is for the statistical significance of the instruments targeting experience, namely the age polynomial, age of youngest and oldest child and dummy for being a mother. All regressions also include all variables in the employment regression model, both instruments and matching variables. Clustered standard errors at the individual level in parentheses under the estimated coefficients.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Years of part-time work experience among women by education, linear regression model

	Education		
	GCSEs	A-levels	Degree
Res. simulated inc. if $h = FT$	-0.014*** (0.004)	-0.013*** (0.004)	-0.011** (0.005)
Diff in res. sim. inc. (FT-PT)	0.006 (0.005)	0.004 (0.006)	0.012* (0.007)
age	-2.932 (6.238)	7.206 (6.016)	6.920 (7.298)
age squared	0.615 (1.796)	-2.301 (1.711)	-1.822 (2.108)
age cubic	0.028 (0.167)	0.278* (0.158)	0.173 (0.201)
age youngest child	0.013 (0.059)	-0.076 (0.072)	0.234* (0.131)
age oldest child	0.162*** (0.055)	0.220*** (0.065)	0.079 (0.114)
mother	-1.316*** (0.408)	-0.233 (0.530)	-2.573*** (0.831)
F -test statistic	113.094	46.957	11.637
p-val	0.000	0.000	0.000
N	9,481	6,845	3,212

Notes: Estimates on BHPS data for the 1991-2008 period. The F statistic is for the statistical significance of the instruments targeting experience, namely the age polynomial, age of youngest and oldest child and dummy for being a mother. All regressions also include all variables in the employment regression model, both instruments and matching variables. Clustered standard errors at the individual level in parentheses under the estimated coefficients. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.